USPTO

STIC Search Report
EIC 2100

STIC Database Tracking Number: 180231

TO: Mohammad A Siddiqi
Location: RND 4C21
Art Unit: 2154
Thursday, February 23, 2006

Case Serial Number: 09/692351

From: Emory Damron
Location: EIC 2100
RND 4B19
Phone: 571-272-3520

Emory.Damron@uspto.gov

Search Notes

Dear Mohammad,

Please find below the rest of your NPL fast and focused search.
References of potential pertinence have been tagged, but please review all the packets in case you like something I didn't.
Please see the email I sent you with regard to this search.

Please contact me if I can refocus or expand any aspect of this case, and please take a moment to provide any feedback (on the form provided) so EIC 2100 may better serve your needs. Good Luck!

Sincerely,
Emory Damron
Technical Information Specialist
EIC 2100, US Patent & Trademark Office
Phone: (571) 272-3520
Emory.damron@uspto.gov

Scientific & Technical Information Center
S T I C
USPTO

STIC EIC 2100 Search Request Form

180231

92

Today's Date: 02/22/2006	**What date would you like to use to limit the search?** Priority Date: 10/19/2000 Other:

Name MOHAMMAD SIDDIQI	**Format for Search Results (Circle One):**
AU 2154 Examiner # 79997	PAPER DISK EMAIL
Room # 4C21 Phone 23976	**Where have you searched so far?** USP DWPI EPO JPO ACM IBM TDB
Serial # 09/692351	IEEE INSPEC SPI Other ________

Is this a "Fast & Focused" Search Request? (Circle One) YES NO
A "Fast & Focused" Search is completed in 2-3 hours (maximum). The search must be on a very specific topic and meet certain criteria. The criteria are posted in EIC2100 and on the EIC2100 NPL Web Page at http://ptoweb/patents/stic/stic-tc2100.htm.

What is the topic, novelty, motivation, utility, or other specific details defining the desired focus of this search? Please include the concepts, synonyms, keywords, acronyms, definitions, strategies, and anything else that helps to describe the topic. Please attach a copy of the abstract, background, brief summary, pertinent claims and any citations of relevant art you have found.

assigning bits to Loca Identifier and assignig Local Identifier to a port in a end node; Multiple Partition environm

I would search in infiniband Network environment the infiniband specification before the filing date will be useful

IBM

606F
TO1

STIC Searcher Emon Damra Phone 23520
Date picked up 2/22/06 Date Completed 2/23/06

sira
Search and Information Resources Administration

XCOPY

STIC Search Results Feedback Form

USPTO

Scientific & Technical Information Center

S T I C

EIC 2100

Questions about the scope or the results of the search? Contact ***the EIC searcher or contact:***

Anne Hendrickson, EIC 2100 Team Leader
272-3490, RND 4B28

Voluntary Results Feedback Form

- *I am an examiner in Workgroup:* 2454 *Example: 2133*
- *Relevant prior art **found**, search results used as follows:*
  - ☐ 102 rejection
  - ☐ 103 rejection
  - ☐ Cited as being of interest.
  - ☐ Helped examiner better understand the invention.
  - ☐ Helped examiner better understand the state of the art in their technology.

  Types of relevant prior art found:
  - ☐ Foreign Patent(s)
  - ☐ Non-Patent Literature
    (journal articles, conference proceedings, new product announcements etc.)
- *Relevant prior art **not found:***
  - ☐ Results verified the lack of relevant prior art (helped determine patentability).
  - ☐ Results were not useful in determining patentability or understanding the invention.

Comments:

Drop off or send completed forms to STIC/EIC2100 RND, 4B28

sira
Search and Information Resources Administration

```
? d s
Set      Items   Description
S1        1031   S INFINIBAND? OR INFINI()BAND?
S2         192   S PLURAL? OR MULTIT? OR MULTIP? OR NUMEROUS? OR SEVERAL? OR MANY OR
MORE(2W)ONE OR ARRAY?
S3          67   S ARRAY? OR ASSORTMENT? OR GROUP? OR COLLECTION? OR ASSEMBLY?
S4          65   S PARTITION? OR PART? ? OR PARTIAL? OR SEGMENT? OR DIVISION?
S5           7   S DEMARCAT? OR SECTION? OR SECTOR? OR PORTION? OR APPORTION?
S6         136   S PARCEL? OR PIECE? OR CHUNK? OR FRACTION? OR AREA? ? OR ZONE? ?
S7         151   S INDEPEND? OR INDIVIDUAL? OR SINGLE? OR SINGULAR? OR DISTINCT? OR
SPECIFIC? OR PARTICULAR?
S8          17   S CERTAIN? OR DEFINIT? OR PRECIS?
S9           1   S ENDNODE? OR END()NODE?
S10         27   S ASSIGN? OR DESIGNAT? OR ALLOCAT? OR EARMARK? OR CONSIGN? OR ALLOT? OR
APPOINT?
S11        125   S DISTRIBUT? OR APPORTION? OR PARCEL? OR PARTIT? OR DISBURS? OR DIVVY?
S12         86   S BIT OR BITS OR PACKET? OR DATAPACKET?
S13          2   S (LID OR LIDS)(7N)(LOCAL? OR IDENT? OR ADDRESS?) OR
LOCAL?()IDENTIF?()ADDRESS? OR LOCAL?()IDENTIF?
S14          0   S (LMC OR LMCS)(7N)(LOCAT? OR MASK? OR CONTROL?) OR
LOCAT?()MASK?()CONTROL?
S15          8   S (CHANNEL? OR ADAPT?)(3N)PORT?
S16         94   S S1 AND S12:S15
S17        146   S S1 AND S10:S11
S18        458   S S1 AND S2:S15
S19        458   S S16:S18
S20         39   S S19 AND PY<2001
S21         45   S S19 NOT PY>2000
S22         45   S S20 OR S21
S23         45   RD  (unique items)
 ; show files
```

[File 2] **INSPEC** 1898-2006/Feb W2

**File 2: Archive data back to 1898 has been added to File 2.*

[File 6] **NTIS** 1964-2006/Feb W1

[File 8] **Ei Compendex(R)** 1970-2006/Feb W2

[File 34] **SciSearch(R) Cited Ref Sci** 1990-2006/Feb W2

[File 35] **Dissertation Abs Online** 1861-2006/Jan

[File 65] **Inside Conferences** 1993-2006/Feb W3

[File 94] **JICST-EPlus** 1985-2006/Nov W4

[File 99] **Wilson Appl. Sci & Tech Abs** 1983-2006/Jan

[File 111] **TGG Natl.Newspaper Index(SM)** 1979-2006/Feb 15

[File 144] **Pascal** 1973-2006/Jan W5

[File 239] **Mathsci** 1940-2006/Mar

[File 256] **TecInfoSource** 82-2006/Feb


```
 ; logoff hold
Holding session beginning:  2/23/2006 8:24:05 AM
Just enter a command to reestablish your session

?  logoff
Ended session:  2/23/2006 8:24:18 AM

? Your settings remain unchanged
?
```

23/5,K/8 (Item 3 from file: 99) Links
·Fulltext available through: USPTO Full Text Retrieval Options

2191097 **H.W. Wilson Record Number:** BAST00065150
InfiniBand promises greater speed, scalability for servers and clusters

Grossman, Steve ;
Electronic Design v. 48 no20 (Oct. 2 2000) p. 25-6
Document Type: Feature Article **ISSN:** 0013-4872 **Language:** English **Record Status:** Corrected or revised record

Abstract: **InfiniBand**, a new I/O switch-fabric architecture, provides greater bandwidth than existing bus-structured architectures. The principal components of an **InfiniBand** network are the switches, the host channel adapters, and the target channel adapters. The network design **segments** 2 Gbyte messages into 256- to 4,096-byte packages and reduces redundancy and hot-pluggability problems. The **InfiniBand** Trade Association has employed a design that leaves room for application vendors to innovate.

Descriptors: Fieldbus (Transmission lines); Computer network architecture ;
InfiniBand promises greater speed, scalability for servers and clusters

Abstract: InfiniBand, a new I/O switch-fabric architecture, provides greater bandwidth than existing bus-structured architectures. The principal components of an InfiniBand network are the switches, the host channel adapters, and the target channel adapters. The network design segments 2 Gbyte messages into 256- to 4,096-byte packages and reduces redundancy and hot-pluggability problems. The InfiniBand Trade Association has employed a design that leaves room for application vendors to innovate.
2000

InfinBand promises greater speed, scalability for servers and clusters
Steve Grossman
Electronic Design; Oct 2, 2000; 48, 20; Career and Technical Education
pg. 25

InfiniBand Promises Greater Speed, Scalability For Servers And Clusters

If bus architectures such as PCI and PCI-X could deliver the higher I/O clock rates and scalability essential for 21st century networks, an I/O architecture known as InfiniBand never would have been hatched. In the early 1990s, 66-MHz processors and 10-Mbit/s networks were considered state-of-the-art. Servers back then cranked out 54 transactions/min., a fairly paltry figure by today's standards. PCI and PCI-X were modeled to meet the needs of the the last decade. With their initial 133-Mbit/s bandwidth and subsequent increases to 1.066 Gbits/s, they did.

Look at where we are now. This March, Intel announced a Pentium III Xeon processor that runs at 1 GHz. As a result, servers have sped past 135,000 transactions/s. It's no wonder that an up to date I/O structure has become necessary. This has given rise to a new I/O structure, InfiniBand, whose rollout is under way. Though still a fledgling specification, version 1.0 should launch this month.

InfiniBand is a network approach, rather than a bus approach, to I/O architecture *(see the figure)*. The key components in an InfiniBand network are the switches, the host channel adapters (HCAs), and the target-channel adapters (TGAs).

The switch is a relatively simple device. It forwards dual-simplex, 8-bit/10-bit encoded packets. These packets are based on two fields that they contain, known as a destination/local ID and a service-level field. Messages of up to 2 Gbytes are segmented into packet lengths ranging from 256 to 4096 bytes, depending upon the application.

If a reliable protocol is selected, it can be guaranteed that every packet will be delivered once and only once, in order,

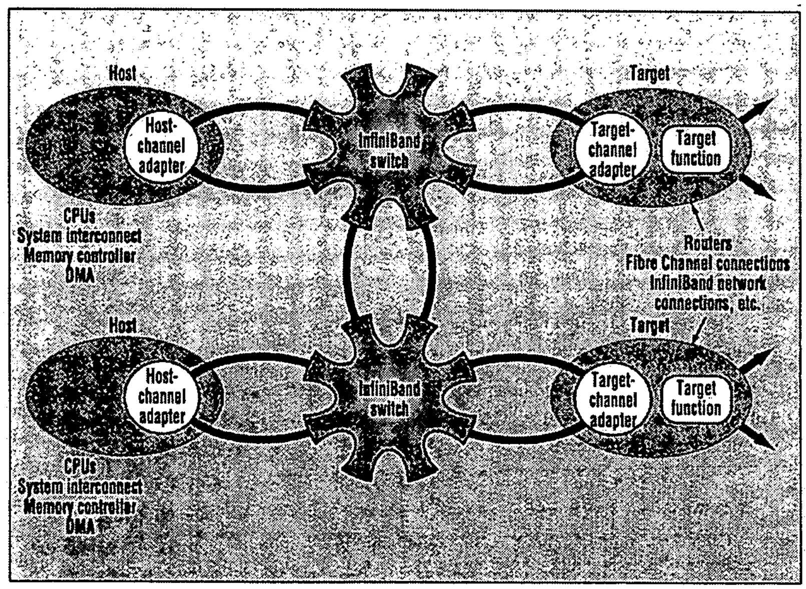

InfiniBand, a switch-fabric architecture, offers far greater bandwidth than the bus-structured architectures of the past. It also eases redundancy and hot-pluggability problems. This network approach segments 2-Gbyte messages into 256- to 4096-byte packets.

uncorrupted. Users are notified if that's not possible. Once packets are sent out, they're reassembled at the far end to complete the transaction.

Aggregate bandwidths are 500 Mbits/s, 2 Gbits/s, and 6 Gbits/s with a 2.5-Gbit/s signaling rate. The InfiniBand Trade Association (IBTA) hopes to eventually boost performance to a higher signaling rate. Bandwidths are scaled depending on how links are aggregated—by 1, by 4, or by 12. Serial links are traditionally described in bits/s, rather than in bytes/s, as is the case with parallel links.

Each InfiniBand width drives 2.5 Gbits/s (250 Mbits/s) in each direction. A 4-wide architecture is 10 Gbits/s (1 Gbit/s per direction). A 12-wide architecture is 30 Gbits/s (3 Gbits/s per direction). From a physical perspective, links may be copper or optical. They will be able to able to drive 20 in. of printed wiring, or 17 m of copper cable, while maintaining a worst-case bit error rate of at least 1×10^{-12}.

Room For Innovation

On the software side, the IBTA wanted to leave a lot of room for applications vendors to innovate. Instead of defining an absolute applications programming interface (API), the IBTA created an abstraction called "Verbs." This innovation defines the functionality that an HCA has to provide. Application vendors, then, know what services are going to be supported. But they're still free to develop individual interfaces, optimizing them for a particular operating system.

From a management standpoint,

one node/switch must emerge as a subnet manager. It can reside in a node or in an HCA. Or, it may be integrated as part of a switch. The subnet manager is responsible for assuring conductivity throughout the fabric. It does this by sending management datagrams. Every InfiniBand device that participates on the fabric has a subnet management agent.

Also, InfiniBand supports unannounced hot-swapping. Designers can just walk up to a module and pull it out. The subnet manager will automatically detect this event.

The IBTA, comprising over 180 companies, came into being in August 1999 as a confluence of two earlier groups: the Next Generation I/O (NGIO) led by Intel, and the Future I/O led by IBM, Compaq, and Hewlett-Packard.

The problems that brought InfiniBand into being are based heavily on the requirements of servers and clusters of servers, sometimes dubbed "farms." Bus architectures lack sufficient headroom. Their capabilities are strained by the voracious demands for data transfers, and in particular, on the Internet and the higher I/O data rates required.

"There is a real crunch at the data centers," says Jean S. Bozman, research director, Commercial Systems and Servers, International Data Corp., Framingham, Mass. Bozman spoke at August's Intel Development Forum in San Jose, Calif.

"A high-speed interconnect such as Infiniband is going to promote flexibility in computer system design. When we have these new, faster links we will be free to move the server pieces farther apart—or arrange them in a little different way. Whereas before it has all been in the confines of a single cabinet or box," according to Bozman. "It will also put an end to fork-lift upgrades," she adds, referring to the practice of removing and replacing servers, en masse, rather than upgrading existing servers. "In fact, expandable servers will enable capacity upgrades on-the-fly."

OEMs looking to participate in the development of InfiniBand-based products have a number of opportunities. Bozman advises vendors to identify early on specific market segments that they believe will adopt InfiniBand. Then, they need to develop plans for phased InfiniBand rollouts by working with software vendors to make sure key applications use InfiniBand APIs.

"Building the InfiniBand infrastructure is going to be kind of a layered approach with, at first, a lot of the technology coming in at the edges," Bozman predicts. She sees InfiniBand arriving in concert with the move in servers from 32-bit computing to 64-bit computing, pointing out that the 64-bit versions will support both 64-bit as well as the large inventory of 32-bit applications now in place.

Early InfiniBand components will most likely be bridge chips and add-on cards for connecting with existing products via existing I/Os. Enlarged support for clustering and server farms will arrive in 2001, tying together existing systems with InfiniBand-type clustering. Full-blown symmetrical multiprocessor servers (SMPs) embodying InfiniBand will probably emerge further down the line.

Some issues remain to be answered, though. For example, it's unclear if InfiniBand will complement or compete with bus architectures, such as the well-entrenched PCI and its successor, PCI-X, which is less than a year old.

For more details, go to the group's web site at *www.infinibandta.org*.

Steve Grossman

Updated Tool Does Kernel-Level Debugging For Real-Time Linux

The Linux Trace Toolkit (LTT) now supports kernel-level debugging. It has been available for application-level debugging for some time, but real-time developers had to contend with basic debugging tools. This version of LTT also supports the Linux hard real-time application interface (RTAI).

LTT is an open-source project supported by Opersys Inc. of Montreal, Canada, and Lineo Inc. of Lindon, Utah. It's distributed under the GNU General Public License, making it freely available to developers. As a graphical tool, it dynamically displays system performance information (*see the figure*). It can be used to determine what process was accessing hardware in a particular time slice. Also, it can highlight I/O device-driver latencies or application dependence on device drivers. It's especially useful in analyzing syncrhonization problems.

Converting regular Linux applications to hard real-time applications under RTAI requires one call, rt_make_hard_real_time(). But making sure the application does what it's intended to do under these constraints is now much easier. LTT support currently works in a single-processor environment. The support hooks have not been added in the multiprocessor Linux kernel. This is a limitation with RTAI, LTT, and the multiprocessing support. LTT already works in this environment for regular applications.

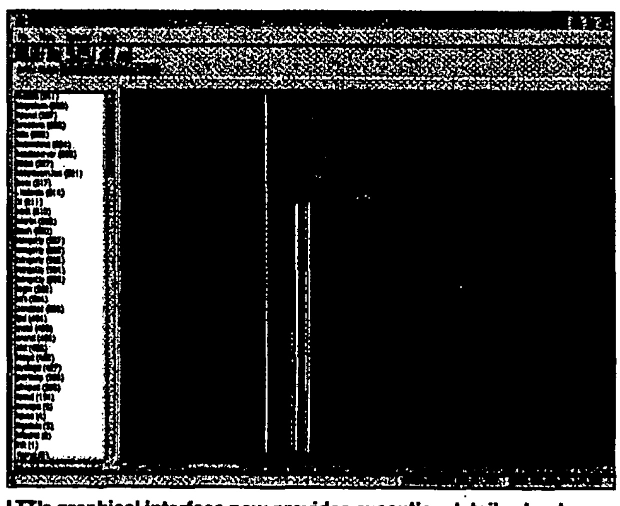

LTT's graphical interface now provides execution details about kernel applications, as shown in this event-flow graph.

A minor change to the LTT trace-file format makes it incompatible with prior versions, though. On the plus side, the binary format is smaller, which is very handy because tracing over long periods of time can generate rather large files. The binary encoding also improves performance.

For more information about the LTT, visit *www.opersys.com/LTT/*.

William Wong